

Attorney Docket No. 19463-0010
Application No. 10/620,150
In reply to Office Action dated June 28, 2005

Amendments to the Drawings:

The attached two sheets of drawings includes changes to Fig. 2 and 4. These sheet replaces the original sheets including Figs. 2 and 4. In Figure 2, previously omitted elements Genlock, and Master timing device, and slave timing devices have been added. In Figure 4, previously omitted element load balancing ratio has been added.

Attachment: Two Replacement Sheets
Two Annotated Sheets Showing Changes

REMARKS/ARGUMENTS

Following this Amendment, independent claims 1, 65, 66, and 67 remain in the present application, with claims 2-32 and 68-76 depending therefrom. Applicants believe that the current amendment does not introduce any new matter and, therefore, respectfully request entry of the present amendment. It is believed that all grounds for objection/rejection have been addressed through the amendment and the following remarks, and that the application is now in condition for allowance. Applicants kindly solicit reconsideration the present invention as currently claimed and allowance the claims.

Previous Restriction Requirement

Claims 33-64 have been withdrawn as the result of an earlier restriction requirement. In view of the examiner's earlier restriction requirement, applicants retain the right to present claims 33-64 in a divisional application.

Objection to the Drawings under 37 CFR §1.83

In response to the objection to the drawings under 37 CFR §1.83 (a), Figures 2 and 4 have been amended to include previously omitted elements. Specifically, Figure 2 has been amended to include a Genlock component within the Synchronization element 232 and to include master and slave timing devices incorporated into the video cards 218, 220, as supported in Paragraph 46 of the printed patent application. Specifically, Paragraph 46 discloses:

The multiple video cards are **synchronized 232** to ensure that all of the video cards in the system are generating video data for the same pixel at the same time. There are multiple methods, well known to those skilled in the art, of achieving this type of synchronization, a detailed discussion of which is beyond the scope of this patent. By way of example, one way to achieve synchronization is by using a **genlock (short for generator locking) mechanism**. A genlock mechanism, generally speaking, synchronizes multiple devices to a specific timing signal. Another method for achieving synchronization between the multiple video cards is to designate the timing regulating device in one of the video cards as a **master timing regulating device** and modify the circuit in the other cards so that the timing regulating devices in those cards act as **slaves** of the master timing regulating devices.

Likewise, Figure 4 has been amended to include a Load Balancing Ratio in the Data Bus 110 as described in Paragraph 53-54 of the printed application. Specifically, Paragraph 53 discloses, “[T]he ‘**load balancing ratio**,’ is assigned through software and, optionally, through user input, and is obtained by the microcontroller 316 **from the computer’s data bus 110**. The load balancing ratio is, in turn, obtained by the video switch controller 320 from the VMH microcontroller 316.” Applicants believe that no new subject matter is being added through these additions which are fully supported in the specification, as described above, and respectfully request entry of the amended figures. Applicants believe that the objection to the Drawings under 37 CFR §1.83(a) has been addressed through this amendment.

Claim Rejection Under 35 USC §112, First Paragraph

The Office action rejected claims 1-32, 65, and 75-76 under 35 USC §112, First Paragraph for containing elements that were not described sufficiently in the specification to enable one of ordinary skill in the art of video graphics processing to make or use the claimed invention. Applicants have carefully reviewed the claims and respectfully submit that the claim elements are sufficiently described in the specification or otherwise known in the field of video graphics processing to allow someone of ordinary skill in the field to practice the invention as claimed. The claims terms raised in the Office Action are known addressed individually.

The Office Action rejected claim 1-3 and 65 that the term “graphics command replicator” is patently indefinite. Applicants respectfully disagree. Specifically, the graphics command replicator, or GCR, is depicted as element 204 in FIG. 2 and the accompanying text, for example in paragraph 39 and 41-42 of the printed publication. As described in paragraph 39:

[T]he GCR module 204 is a software program that resides between the computer application and multiple instances of the API module 203,205. The GCR identifies and intercepts API commands 202 issued by the application 200 before those commands reach the API module instances 203,205. Once intercepted, the GCR module 204 generates multiple, modified API command streams 206,208.... The number of modified API command streams 206,208, and of instances of the API module 203,205, in this case two (2), is equal to the number of video cards being employed in the system.

Applicants thus believe that it would be clear to someone of ordinary skill in the art that the graphics command replicator is a software program to receive a API commands from an application and replicates these received commands for each video card. Thus, Applicants suggest that this claim term is described sufficiently in the specification to enable one of ordinary skill in the art of video graphics processing to make or use the claimed invention.

The Office Action further rejected claim 1 for containing the term “off-the-shelf video cards,” and Applicants have removed this phrase to expedite examination even though Applicants believe that someone of ordinary skill in the field of video processing would be familiar with the term.

Regarding the rejection of claim 20 for the use of the term “Genlock”, Applicants respectfully suggest that a Genlock is well known in the field of computer science. For example, one online technical encyclopedia defines “Genlock” (for Generator Lock) as: “is a technique where the output of one system is used to synchronize another. Genlock is a common technique in video and digital audio engineering.” See <http://en.wikipedia.org/wiki/Genlock>. Applicants consequently suggest that “gunlock” would be clear to someone of ordinary skill in computer science, and thus, this claim term is described sufficiently in the specification to enable one of ordinary skill in the art processing to make or use the claimed invention.

As for the rejection of claims 21 and 22 for the use of the terms “master timing devices” and “slaves of said master timing regulating device,” Applicants again respectfully suggest that these terms is well known in the field of computer science. For example, Federal Standard 1037C entitled Telecommunications: Glossary of Telecommunication Terms, a U.S. Federal Standard issued by the General Services Administration pursuant to the Federal Property and Administrative Services Act of 1949, defines a master clock to be a clock that independently produces a timing signal. and a slave clock to be a clock that is coordinated with the master clock. Slave clock coordination is usually achieved by phase-locking the slave clock signal to a signal received from the master clock. To adjust for the transit time of the signal from the master clock to the slave clock, the phase of the slave clock may be adjusted with respect to the signal from the master clock so that both clocks are in phase. Thus, the time markers of both clocks, at the output of the clocks, occur simultaneously. For more information, please see Federal Standard 1037C at <http://www.its.bldrdoc.gov/fs-1037/fs-1037c.htm>. Applicants believe that

master and slave timing devices would be well known to someone of ordinary skill in computer science, and that these claim terms are described sufficiently in the specification to enable one of ordinary skill in the art processing to make or use the claimed invention.

Turning now to the rejection of claims 27-32 and 71-76 under 35 USC §112, paragraph 1 for use of the phrase "Load Balancing Ratio," Applicants again respectfully suggest that the concept of load balancing is well known in the field of computer science, particularly networking. "Load balancing is a technique used to spread work between many processes, computers, disks or other resources" [Http://en.wikipedia.org/wiki/Load_balancing](http://en.wikipedia.org/wiki/Load_balancing). For example, load balancing is commonly used in networking to spread duties and burdens across multiple servers, and the load balancing functionality is typically included in hardware (<http://www.intel.com/design/network/solutions/lbs/>) and software (<http://www.microsoft.com/technet/prodtechnol/windowsserver2003/library/ServerHelp/543cd17d-899a-401e-a01e-3847d8affc52.mspx>). Specifically, referring to Network Load Balancing offered by Microsoft Window Server® described online at <http://www.microsoft.com/technet/prodtechnol/windows2000serv/deploy/confeat/nlbovw.mspx>:

Network Load Balancing scales the performance of a server-based program, such as a Web server, by distributing its client requests among multiple servers within the cluster. With Network Load Balancing, each incoming IP packet is received by each host, but only accepted by the intended recipient. The cluster hosts concurrently respond to different client requests, even multiple requests from the same client. For example, a Web browser may obtain the various images within a single Web page from different hosts in a load-balanced cluster. This speeds up processing and shortens the response time to clients.

Each Network Load Balancing host can specify the load percentage that it will handle, or the load can be equally distributed among all of the hosts. Using these load percentages, each Network Load Balancing server selects and handles a portion of the workload. Clients are statistically distributed among cluster hosts so that each server receives its percentage of incoming requests. This load balance dynamically changes when hosts enter or leave the cluster. In this version, the load balance does not change in response to varying server loads (such as CPU or memory usage). For applications, such as Web servers, which have numerous clients and relatively short-lived client requests, the ability of Network Load Balancing to distribute workload through statistical mapping efficiently balances loads and provides fast response to cluster changes.

Thus, it can be seen that the concept of load balancing ratio (or "Load Percentage" as described above in the Microsoft software solution) is well known in the field of computer science.

Applicants respectfully suggest that "load balancing ratio" would be clear to someone of ordinary skill in computer science, and thus, this claim term is described sufficiently in the specification to enable one of ordinary skill in the art processing to make or use the claimed invention.

Rejection under 35 USC §103

The Office Action rejected the pending claims in the present patent application under 35 USC §103 as being obvious in the view of the combination of Taylor (WO97/14133) and Callway (EP 1061434). Applicants have carefully reviewed the cited references and believe that the present invention is patentably distinct from these references, either separately or in combination. To expedite examination of the present invention, applicants herein amend the claims to clarify and highlight the patentable distinctions. For example, Claim 1 is amended to recite:

An accelerated graphics processing subsystem comprising: a graphics command replicator; a plurality of video cards; a mechanism to synchronize the signal by said plurality of video cards; and a video merger hub,

wherein said routing switch triggering intervals are determined on the basis of: the vertical refresh rate of said output signals from said plurality of video cards; the vertical resolution of said output signals from said plurality of video cards; and the load balancing ratio assigned to each of said plurality of video cards.

In summary, the invention provides a graphics processing subsystem used in personal computers to accelerate graphics processing for multiple applications, such as gaming, video and photographic editing, multimedia displays, computer graphics imagery (CGI), etc. The invention specifically provides a solution to coordinate the operation of multiple video cards but dividing a display screen into multiple areas and then assigning a separate video card to each of the display areas. The present invention then further optimizes the performance of the graphics subsystem by using a load balancing ratio to allocate graphics processing between the multiple video cards. As disclosed in the specification, the load balancing ratio may be dynamically determined according to the measured performance of the video processing as defined by the previous processing tasks and overall capability. In this way, the video cards may all operate at close to

optimal efficiency while minimizing situations where one or more of the cards are waiting idly while waiting for another video card to complete its assigned processing tasks.

Referring now to FIG. 2 of the present application, the present invention includes, a Graphics command replicator (GCR) module 204 that is a software program that resides between the computer application 200 and multiple instances of the API module 203,205. The GCR identifies and intercepts API commands 202 issued by the application 200 before those commands reach the API module instances 203,205. Once intercepted, the GCR module 204 generates multiple, modified API command streams 206,208. The modified API command streams 206, 208 are received by the API module instances 203,205 which in turn each generate a command stream 207,209 that is received and processed by its assigned video card driver 210,212 and their respective video cards 218,220. The number of modified API command streams 206,208 and of instances of the API module 203,205, in this case two (2), is equal to the number of video cards 218,220 being employed in the system. The API streams are generated in such a way that each of the video cards 218,220 will generate only the pixels that are contained within a particular region of the screen assigned to that video card. Each of the command streams 207,209 is then processed by its assigned video card driver 210,212, which in turn, issues GPU commands 214,216 to a respective video card 218,220. Each video card 218,220 in turn generates a video signal 222,224 corresponding to its respective portion of the screen. A synch device 232 coordinates the operation of the video cards 218,220 using known techniques such as a genlock or master/slave timing regulating devices.

The video data output from the multiple video cards 218, 220 is synchronized and combined using a video merger hub (VMH) 226 generally depicted in FIG. 4. Specifically, the principal components of the VMH are a video switch 322, a video switch controller 320, a microcontroller 316, and a video output 330. Typically, each video signal received by the VMH is composed of a video data component 308,310 and a synchronization component 312,314, where the video data component 308,310 is comprised of red, green and blue ("RGB") values for the pixel that is being drawn at a particular time and the synchronization component 312,314 includes vertical and horizontal synchronization signals which determine the vertical and horizontal position (i.e., coordinates) of the pixel that is being drawn at a particular time. As the video signals 222,224 arrive at the VMH, their video data components 308,310 are routed to the

video switch 322. The video switch 322 is, in turn, controlled by the video switch controller 320 which receives the synchronization components 312,314. During each screen refresh cycle, the video switch intelligently and sequentially routes the video data component from the various video signals 222,224 in such a manner that a single, seamless combined video signal 228 is then transferred from the video output 330 of the VMH to the display device 168 along with the synchronization components 312,314 which essentially "pass through" the video switch controller. The video switch controller 320 uses the vertical refresh rate, vertical resolution and a load balancing ratio to calculate the triggering intervals to be used to generate the combined video signal 228. As described above, the "load balancing ratio" represents the percentage of the screen to be assigned to each video card 218,220. A simple test feedback loop program may be used to dynamically adjust the load balancing ratio based on the load of each of the video cards, on a dynamic or frame by frame basis, can maximize the throughput of the combined GPUs.

Applicants respectfully submit that the invention is distinguishable over the applied references for at least the reasons described below.

As an initial observation, Applicants submit that Taylor does not teach or suggest the present invention. Referring now FIG. 1 (reproduced below), Taylor is a patent application in which multiple display controllers, or video cards, 104, each drive a separate section of a single display 110. As described in the first paragraph of the Invention Summary starting at line 27 of page 3, Taylor specifically addresses the question of how to separately address each of the video cards 104, and proposes to assign unique bus identifiers, or addresses, to each of the video cards 104. The program data are processed to determine the relevant display area and an associated video card 104. The graphics commands are then appended with prefix addresses for the associated video card 104 and sent over a system bus for receipt from the designed video card having the appended address. While Taylor describes the use of a digital to analog converter (DAC)/palette (display driver) 109 that receives digital data from each of the controllers 104, there is no discussion in Taylor of how the display driver 109 coordinates the pixel data from each of the display controllers 104.

Consequently, Callway states in Paragraphs 4-9 at page 2 that it is known to use multiple graphics processor by partitioning a display area and associated a processor to each of the processors, but that this technique had several problems processing coordination, data allocation,

and load balancing, likely through the lack of processor synchronization coordination as described above in the discussion of Taylor. As described in the Abstract, Callway attempts to address these identified problems by coordinating the video graphics adaptors (VGAs) 110, 120 by a controller 130 to render video data to a common port 151. **“Each VGA 110, 120 will render an entire frame of video and provide it to the output port 151 through a switch 1.** The next adjacent frame will be calculated by a separate VGA and provided to an output port through the switch.” As described in paragraphs 21-22 of Callway, the controller 130 sends out timing signals to coordinate and synchronize the VGAs.

Thus, it can be seen that Callway coordinates the video cards by having the cards draw alternating frames, and thus specifically teaches away from the display division techniques used in the present application. Furthermore, it would seem that Callway is generally incompatible with Taylor for similar reasons, and consequently, combination of those two references seems improper.

In conclusion, it can be seen that Taylor does not use any type of coordination or central control of the multiple video cards and Callway uses frame-by-frame labor coordination (time coordination), and not the display area-based coordination of the present invention.

Furthermore, Talyor and Callway cannot be used to anticipate and make obvious the present invention's use of a controller to the coordinate the different pixel locations from the outputs of multiple video cards to render a single display using the vertical refresh rate of the output signals from the video cards; the vertical resolution of the output signals from the video cards; and the load balancing ratio assigned to each of the video cards. Simply, Taylor and/or Callway do not suggest the use of a load balancing ratio to dynamically adjust video card assignments as needed to efficient processing load allocation. Consequently, in Callway or Taylor, a relatively fast display card may sit idle after processing the next frame while waiting for the other card to complete its processing of the current frame.

Thus it can be seen that the combination of Taylor and Callway does not anticipate or make obvious the present invention as recited in claim 1 for at least the reason that the reference do not teach for the coordination of multiple video cards through display screen division techniques and that the references do not suggest the use of a load balancing between the separate video cards. Likewise, the other remaining independent claims 65-67 are allowable over the

combination of Taylor and Callway on similar grounds. The remaining dependent claims 2, 5-22, 28-32, and 72-76 are similarly allowable as depending from allowable claims.

Applicants further believe that the combination of Taylor and Callway does not suggest the present invention as provided in some of the dependent claims. For example, Taylor and Callway do not teach or suggest:

- As specified in claim 2, graphics command replicator is a software module that intercepts graphics commands issued by an application and generates multiple, modified graphics command streams; the number of said multiple, modified graphics command streams is equal to the number of said plurality of video cards; and each of said multiple, modified graphics command streams is received by a separate video card selected from said plurality of video cards.
- As specified in claim 6, the intercepted graphics commands are API commands and said multiple, modified graphics command streams are multiple, modified API command streams.
- As specified in claim 7 the multiple, modified API command streams are each received by a separate instance of an API module and wherein each of said API module instances generates a command stream which is processed by a separate video card selected from said plurality of video cards.
- As specified in claim 8, the aid modifications to said multiple, modified graphics command streams are accomplished by incorporating therein a clipping command.
- As specified in claim 9, the clipping command is a 2D clipping command.
- As specified in claim 10, the clipping command is a 3D clipping command.
- As specified in claim Claim 17, each of said plurality of video cards is equipped with a single GPU.

- As specified in claim 18, each of said plurality of video cards is equipped with a plurality of GPUs.
- As specified in claim 19, the plurality of video cards is comprised of a combination of video cards equipped with a plurality of GPUs and video cards equipped with a single GPU.
- As specified in claim 20, the mechanism to synchronize the signal output by said plurality of video cards is a genlock mechanism.
- As specified in claim 21, the mechanism to synchronize the signal output by said plurality of video cards consists of designating the timing regulating device in one of said plurality of video cards as a master timing regulating device and designating timing regulating devices in the remainder of said plurality of video cards as slaves of said master timing regulating device.
- As specified in claim 22, the timing reference sources for said master and slave timing regulating devices are timing reference sources selected from the group consisting of piezoelectric crystals, programmable crystals, oscillators, programmable oscillators and combinations thereof.
- As specified in claim 28 and 72, said load balancing ratio is transmitted by said microcontroller to said video switch controller.
- As specified in claim 29 and 73, said load balancing ratio is assigned by a user through software.
- As specified in claim 30 and 74, said load balancing ratio is equal for each of said video cards.

Attorney Docket No. 19463-0010

Application No. 10/620,150

In reply to Office Action dated June 28, 2005

- As specified in claim 31 and 75, said load balancing ratio is based on each of said video cards' graphics throughput.
- As specified in claim 32 and 76, said load balancing ratio is dynamically adjusted to maximize the throughput of said subsystem by utilizing a test feedback loop program which measures the load on each of said video cards and makes appropriate adjustments.

Attorney Docket No. 19463-0010
Application No. 10/620,150
In reply to Office Action dated June 28, 2005

Conclusion


In view of the foregoing, the Applicants respectfully request that the Examiner considers the above-noted amendment when the application is examined on its merits and the timely allowance of the pending claims. The Examiner is invited to contact Applicants' undersigned representative to expedite prosecution.

If there are any fees due in connection with the filing of this response, please charge the fees to our Deposit Account No. 50-1349.

Respectfully submitted,

Dated: December 28, 2005

HOGAN & HARTSON LLP
555 13th Street, N.W.
Washington, D.C. 20004
Telephone: 202-637-5564
Facsimile: 202-637-5910
e-mail: ajvaidya@hhlaw.com
Customer No. 24633

By: 
Ajit J. Vaidya
Registration No. 43,214

David D. Nelson
Registration No. 47,818




Fig. 2




Fig. 4